The Stanford Packet Radio Network*

Michael J. Flynn, Curtis Spangler, and Andrew Zimmerman

Computer Systems Laboratory, Electrical Engineering Department

Stanford University, Stanford CA 94305-2192

Abstract

The Stanford Packet Radio Network is designed to provide from 9600 to about 56,000 baud access to "line of sight" users (within about 30 miles) of SUN—the Stanford University Network. The cost objective is $1,000 per user for the marginal cost of modem, transceiver, etc. Basic testing has been completed, and a preliminary network will begin operations in early 1986.

Introduction

The character of the computing system has been changing in recent years, characterized by a movement away from simple time-sharing systems supporting users with simple terminal facilities, towards a computational network system.

In computational networks the user has a workstation as a terminal. This workstation becomes the primary vehicle for editing, preliminary computation, and storage. The seemingly subtle difference in upgrading a "dumb" terminal to a "smart" terminal radically alters the network support requirements. The network must provide access to common data files, specialized resources (laser-printers, typesetting equipment, etc.), as well as high-speed computational engines, software resources such as compilers, etc. Also, the nature of the network interaction is quite different from the earlier time-shared model. In a time-sharing system the user sits at a terminal sending a string of characters to a single processor for processing and storage. The communications channel between the terminal and the central processor system is characterized by relatively long connect times using relatively low bandwidth. Ordinary voice-grade lines supporting 1200 baud are generally adequate for such systems, although the time to fill a screen from processor to terminal is more comfortable for the user at slightly higher speeds. In the computational network model, the workstation itself is the primary user processing and storage facility. After a workstation session, the user prepares a file for expeditious transfer to a central site for further processing. Thus, the need is for high bandwidth and short connect times.

In response to this need, LANs (Local Area Networks) such as Ethernet have provided single-line coax facilities to transmit 10Mb per second in a packet based network. Such technology is perfectly adequate within a range of a thousand feet or so. Beyond that, users must resort to one of a number of WAN (Wide Area Network) techniques. This may consist of simple voice-grade lines with attendant bandwidth limitations, or specially conditioned and dedicated lines to provide higher bandwidth. The ARPAnet is a prime example of a WAN spanning the country and many parts of the globe, providing a 56 Kbaud bandwidth. These dedicated high-bandwidth WANs tend to be very expensive. The motivation for a packet radio network at Stanford is to fill the gap between LANs and WANs with an intermediate low cost, high bandwidth network which we will label LOSAN (Line Of Sight Area Network), spanning distances of up to 50 miles with very low transmission cost. While the network can handle many types of interaction, it is designed to be maximally efficient in serving the computational network model. Short, packet-oriented sessions move files at high speed from a base site at Stanford to a user site for detailed processing. The motivation for providing such a network (Figure 1) is illustrated by some parameters (measured about two years ago) for the movement of a 70,000 byte file on a 1200 baud telephone line.

The protocol between the user and the host was arranged in a typical time-sharing manner, where the file is broken up into 100 character packets and the operating system polls other users between transmission of adjacent packets. Much of the transmission time is simply absorbed in the operating system overhead. A simple provision of a file gateway (a processor that will collect the file and provide dedicated taransmission facilities) provides a significant improvement in performance. Still, 1200 baud is an annoying limitation—the higher the bandwidth, the better the transmission.

The Stanford Packet Radio System

The computing network at Stanford consists of a large number of computers connected through both local and

*Supported in part by NASA under contract NAGW-419. The IBM Corporation has generously provided much of the equipment.

	Data rate	Host operation	Time to complete file transfer
1.	1200 baud	Time-sharing	8 hours
2.	1200 baud	File gateway	45 minutes
3.	9600 baud	"	6 minutes
4.	56,000 baud	"	1 minute

Figure 1: Example File Movement

campus-wide ethernet based technology. With the packet radio system, we use a VAX/750 as a SUN (Stanford University Network) gateway, providing the capability to move large files from any other processor on campus to the gateway processor for subsequent transmission to the remote user. The packet radio system thus has two sites: the base site, which is on campus, and the remote user site. The basic goals for the packet radio network include the following:

- Remote users initially will be constrained to use the IBM PC or a compatible such device.
- The marginal additional cost for a remote user for the radio, modem, etc. is limited to $1,000.
- The host site is expected to be more expensive, perhaps up to $10,000.
- The design target is a 56 Kbaud transmission rate.

System Design and Components

The system is configured as a two-channel star network, with the base site as the center of the star. All transmission must go through the base site; there is no peer-to-peer communication. The system is arranged in a two-channel pair. The downlink provides communication from the base site to the remote users, and the uplink channel provides communication from the users to the base station. The components include:

1. The workstation

The current system is designed about the use of the IBM PC or PC-compatible terminal using an SDLC interface to the modem.

2. The modem

The modem is designed to fit on one board and to plug into the IBM PC, and uses the TI 320 signal processing microprocessor. The modem uses QAM—Quadrature Amplitude Modulation—for modulating data on the data carrier. The QAM technique uses a combination of both phase encoding and amplitude encoding to represent multiple bits on a single sinusoidal cycle. Typically, data states are separated by 45 degrees, providing eight distinct states for a single sinusoidal cycle. If we couple this with two levels of amplitude, we create 16 states or the equivalent of 4 bits of information, which can be represented in a single sinusoidal cycle. We are experimenting with a version of QAM called *polar QAM*, developed by Brett Glass at Stanford. With this technique, phases are shifted by 60 degrees across the cycle format and amplitude is represented by three distinct states—the combination creates six degrees of phase states and three amplitude states for a total of 18 unique data states. This allows the encoding of four data bits plus the availability of two hidden states useful for control signalling from the host to the remote site.

QAM is recursively modulated (imposed) on the FM carrier of the transceiver. The QAM data carrier replaces the microphone input and audio output to the transceiver (the audio filtering can be bypassed to improve data transmision).

3. The remote radio

The transceiver is a small, commercially available 10 watt output receiver-transmitter commonly used by amateurs and by commercial mobile stations. The devices we currently use were bought for under $300 and they operate at about 440 MHz or about an 80 cm wavelength. At these frequencies, transmissions are not much affected by the water vapor content of the atmosphere, and they are less affected than higher frequencies by solid objects. The relatively short wavelength offers some advantages in antenna size.

4. The remote antenna

The remote antenna consists of a six or eight element directional antenna which is aimed at the base site. The high directionality of the antenna compensates for the relatively low output power of the transmitter. The overall size of the antenna is less than a conventional high quality TV antenna.

5. The base station

The design for the base station includes a higher power transmitter (100 watts) in full duplex operation with the receiver. Thus, antennas for both transmitter and receiver must be provided, together with full duplex modem operation. The base station antennas must be omnidirectional or, since by the nature of our geography we aim largely in a pattern to the east of our base station, must span at least a 180 degree forward range.

Software/Network

Software for the packet radio network is broken into three levels: applications, transport protocols, and data link control (see Figure 2). The application software and the transport protocols is provided by the TCP-IP package developed at M.I.T.[1] The data link control is provided as a set of routines (developed at Stanford) to interface between the TCP-IP package and the SDLC hardware in the workstation.

Applications

Applications provided by the software include telnet, tftp, and other miscellaneous servers. Telnet is a virtual

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/standford-packradnet

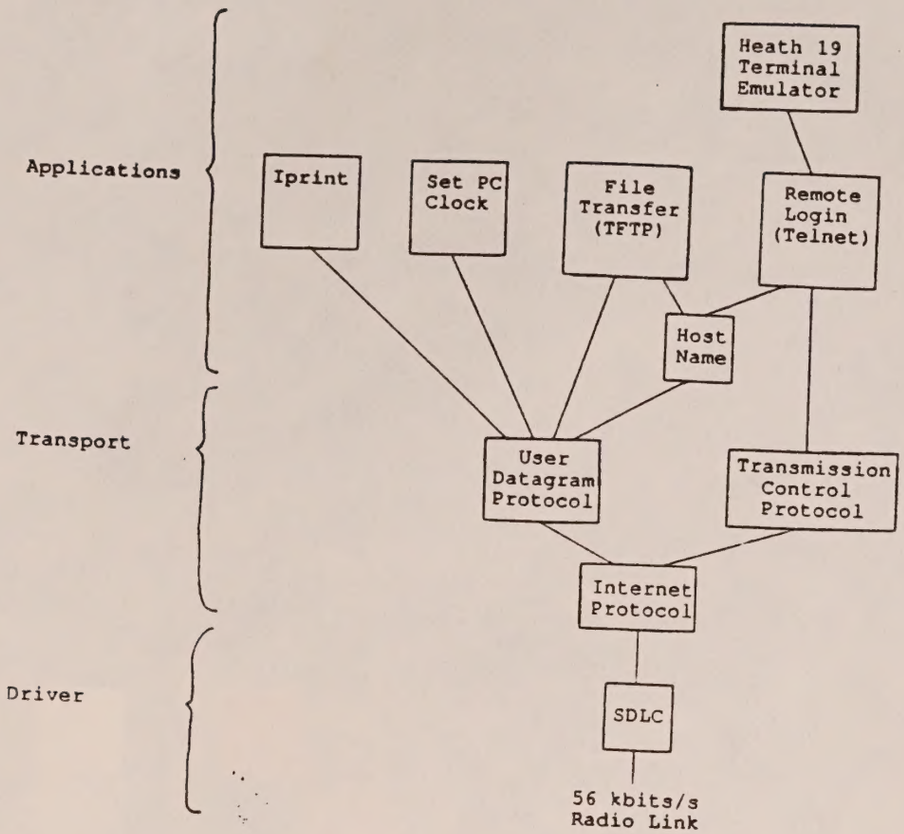

Figure 2: TCP-IP Hierarchy

terminal for remote login to other computers on the network. File transfer is provided by tftp, which permits upload and download of files between the workstation and other hosts on the network. Other useful applications in the TCP-IP package are Iprint (a print server) and setclock (a clock server).

Transport Protocol

The network uses the TCP-IP (transmission control protocol–internet protocol.) This protocol is the same as that used on the Stanford University Network (SUN) ethernet, eliminating the need for the gateway to perform protocol conversion. In this manner, the gateway becomes a simple store and forward node, as well as a concentrator between the packet radio network and the SUN. During transmission between the gateway and the remote node, this protocol is further packed in the information field of the SDLC protocol used as the Data Link.

Data Link

Flow control and error detection over the radio section of the network are implemented using Synchronous Data Link Control (SDLC). This provides a method of retransmission of data when two or more packets collide on the single frequency used by all workstations to communicate to the gateway. SDLC requires that all packets be acknowledged within a given time. If an acknowledgment is not received, the packet is retransmitted. SDLC is also used to synchronize the modem and allows for high speed data rates through the use of direct memory access (DMA.)

Status

We have completed path tests and field studies using standard QAM at 9600 baud. We will initially (by early 1986) offer service to about 10 users (at a time) with a low power (10 watt) base station. We expect to cover only the Palo Alto area during a trial period of about 12 to 18 months, slowly enhancing the system and the service range.

References

[1] J. H. Saltzer, D. D. Clark, J. L. Romkey, and W. L. Gramlich, "The Desktop Computer as a Network Participant," *IEEE Journal on Selected Areas in Communications* Vol. SAC-3, No. 3, May 1985